Primal-dual methods for BSDEs

Christian Bender
(joint work with C. Gärtner, N. Schweizer, J. Zhuo)

Saarland University

FiME Seminar, Paris, Oct. 14h, 2016

Outline

1. Intro to numerical methods for BSDEs
2. Primal-dual methodology for optimal stopping
3. Primal-dual methodology for convex dynamic programs
4. Monte Carlo implementation
5. Numerical example
6. Extensions

Introduction to BSDEs

BSDEs are equations of the form

$$Y_t = \xi - \int_t^T f(s, Y_s, Z_s)ds - \int_t^T Z_s dW_s,$$

where

- W is a D-dimensional Brownian motion
- ξ is the terminal condition ($\mathcal{F}_T$-measurable; filtration generated by W).
- $f : \Omega \times [0, T] \times \mathbb{R} \times \mathbb{R}^D \to \mathbb{R}$.

The solution is a pair (Y, Z) of adapted, square integrable processes.

Intuition: Z_t steers the system Y_t into the terminal condition without making use of future information.

Introduction to BSDEs

Some applications (roughly speaking . . .):

- Mathematical finance: Nonlinear pricing and hedging
  Y_t corresponds to the option price.
  Z_t corresponds to a hedging portfolio.
- Drift control:
  Y_t corresponds to the value process.
  An optimal control can be expressed in terms of (Y, Z).
- Stochastic representations of parabolic PDEs:
  Y_t corresponds to the solution of the PDE.
  Z_t corresponds to the gradient of the solution.

Time discretization

'Backward Euler scheme' (Zhang, 2004; Bouchard and Touzi, 2004)

- Define $h = T/n$ and $t_i = ih, i = 0, 1, \ldots, n$, and $\Delta W_{i+1} = W_{t_{i+1}} - W_{t_i}$ and write formally

$$Y_{t_i} \approx Y_{t_{i+1}} - f(t_i, Y_{t_i}, Z_{t_i})h - Z_{t_i}\Delta W_{i+1}$$

Time discretization

'Backward Euler scheme' (Zhang, 2004; Bouchard and Touzi, 2004)

- Define $h = T/n$ and $t_i = ih$, $i = 0, 1, \ldots, n$, and $\Delta W_{i+1} = W_{t_{i+1}} - W_{t_i}$ and write formally

$$Y_{t_i} \approx Y_{t_{i+1}} - f(t_i, Y_{t_i}, Z_{t_i})h - Z_{t_i}\Delta W_{i+1}$$

- Taking conditional expectation yields

$$Y_{t_i} \approx E_{t_i}[Y_{t_{i+1}}] - f(t_i, Y_{t_i}, Z_{t_i})h$$

Time discretization

'Backward Euler scheme' (Zhang, 2004; Bouchard and Touzi, 2004)

- Define $h = T/n$ and $t_i = ih$, $i = 0, 1, \ldots, n$, and $\Delta W_{i+1} = W_{t_{i+1}} - W_{t_i}$ and write formally

$$Y_{t_i} \approx Y_{t_{i+1}} - f(t_i, Y_{t_i}, Z_{t_i})h - Z_{t_i}\Delta W_{i+1}$$

- Taking conditional expectation yields

$$Y_{t_i} \approx E_{t_i}[Y_{t_{i+1}}] - f(t_i, Y_{t_i}, Z_{t_i})h$$

- Multiplying with ΔW_{i+1} and taking expectation yields

$$Z_{t_i} \approx \frac{1}{h} E_{t_i}\left[Y_{t_{i+1}}\Delta W_{i+1}\right]$$

Time discretization

'Backward Euler scheme' (Zhang, 2004; Bouchard and Touzi, 2004)

- Define $h = T/n$ and $t_i = ih, i = 0, 1, \ldots, n$, and $\Delta W_{i+1} = W_{t_{i+1}} - W_{t_i}$ and write formally

$$Y_{t_i} \Delta W_{i+1} \approx (Y_{t_{i+1}} - f(t_{t_i}, X_{t_i}, Y_{t_i}, Z_{t_i})h)\Delta W_{i+1} - Z_{t_i}(\Delta W_{i+1})^2$$

- Taking conditional expectation yields

$$Y_{t_i} \approx E_{t_i}[Y_{t_{i+1}}] - f(t_i, Y_{t_i}, Z_{t_i})h$$

- Multiplying with ΔW_{i+1} and taking expectation yields

$$Z_{t_i} \approx \frac{1}{h} E_{t_i}\left[Y_{t_{i+1}} \Delta W_{i+1} \right]$$

Time discretization – three variants

'Implicit scheme' (Bouchard, Touzi, 2004)

- Define $h = T/n$ and $t_i = ih, i = 0, 1, \dots, n$, and $\Delta W_{i+1} = W_{t_{i+1}} - W_{t_i}$,

$$\begin{aligned} Y_n^{impl} &:= \xi \\ Y_i^{impl} &:= E_{t_i}[Y_{i+1}^{impl}] - f(t_i, Y_i^{impl}, Z_i^{impl})h \\ Z_i^{impl} &:= \frac{1}{h} E_{t_i}\left[Y_{i+1}^{impl} \Delta W_{i+1}\right] \end{aligned}$$

- Drawback: Y_i^{impl} appears on both sides of the equation.
- Advantage: Good stability properties even when f has polynomial growth in y (Lionnet, dos Reis, Szpruch, 2015).

Time discretization – three variants

'Explicit scheme 1' (Zhang, 2004; Lemor, Gobet, Warin, 2006)

- Define $h = T/n$ and $t_i = ih, i = 0, 1, \ldots, n$, and $\Delta W_{i+1} = W_{t_{i+1}} - W_{t_i}$,

$$\begin{aligned} Y_n^{expl} &:= \xi \\ Y_i^{expl} &:= E_{t_i}[Y_{i+1}^{expl}] - E_{t_i}[f(t_i, Y_{i+1}^{expl}, Z_i^{expl})]h \\ Z_i^{expl} &:= \frac{1}{h} E_{t_i}\left[Y_{i+1}^{expl} \Delta W_{i+1}\right] \end{aligned}$$

Time discretization – three variants

'Explicit scheme 2' (Fahim, Touzi, Warin, 2011)

- Define $h = T/n$ and $t_i = ih, i = 0, 1, \ldots, n$, and $\Delta W_{i+1} = W_{t_{i+1}} - W_{t_i}$,

$$
\begin{aligned}
Y_n^n &:= \xi \\
Y_i^n &:= E_{t_i}[Y_{i+1}^n] - f(t_i, E_i[Y_{i+1}^n], Z_i^n)h \\
Z_i^n &:= \frac{1}{h} E_{t_i}\left[Y_{i+1}^n \Delta W_{i+1}\right]
\end{aligned}
$$

- Under appropriate assumptions on the dependence of ξ and f on ω and a uniform Lipschitz condition on f in (y, z), all three schemes converge at a rate of $1/2$.
- Huge literature on time discretization schemes for BSDEs.

Backward dynamic program – setting

On a filtered probability space in discrete time $(\Omega, \mathcal{F}, \mathcal{F}_i, P)_{i=0,\dots,n}$ we consider the dynamic program

$$Y_i = F(i, E_i[\beta_{i+1} Y_{i+1}]), \quad 0 \le i \le n-1; \quad Y_n = \xi$$

where

- ξ is an $\mathcal{F}_n$-measurable integrable random variable;
- $(\beta_i)_{i=1,\dots,n}$ is an adapted bounded, $\mathbb{R}^{1+D}$-valued process;
- $F(i, \omega, z)$ is Lipschitz in z (with constant independent ω) and $\mathcal{F}_i$-measurable for fixed $z \in \mathbb{R}^{1+D}$.

BSDE setting: $\beta_i^{(0)} = 1$, $\beta_i^{(d)} = [\Delta W_{d,i}]/h$, $d \ge 1$, where $[\Delta W_{d,i+1}]$ are truncated increments of a D-dim. Brownian motion; $F(i, z) = z_0 + f(t_i, z_0, \dots, z_D)h$.

Backward dynamic program

- Typical situation: There is a deterministic function $z(i,\cdot):\mathbb{R}^{\mathcal{D}}\to\mathbb{R}^{1+D}$ and a $\mathbb{R}^{\mathcal{D}}$-valued stochastic process such that
$$E_i[\beta_{i+1}Y_{i+1}]=z(i,X_i).$$
- Example: X_i Markovian, $F(i,\cdot)$ and ξ depend on ω through X_i (resp. X_n) only, β_{i+1} independent of $\mathcal{F}_i$.
- Numerical algorithms typically try to approximate $z(i,x)$.
- Difficulties: High order nesting of conditional expectations within the dynamic program (when n is large); conditional expectation must be replaced by an approximate operator which can be nested without exploding cost.

Nested conditional expectations

- Plain Monte Carlo is not applicable ...

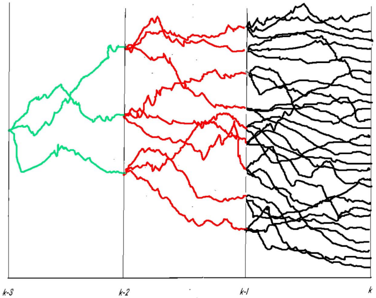

- In the numerical experiments, we apply least-squares Monte Carlo (Longstaff-Schwartz, 2001; Lemor, Gobet, Warin, 2006), but other choices are possible: quantization (Bally, Pages, 2003), Malliavin Monte Carlo (Bouchard, Touzi, 2004), cubature on Wiener space (Crisan, Manolarakis, 2012), ...

Conditional expectations via least squares Monte Carlo

Pseudo-Algorithm

1. Choose $(1+D)$ row vectors of basis functions

$$\psi_d(i,x) = (\psi_{d,1}(i,x),\ldots,\psi_{d,K}(i,x));\ x \in \mathbb{R}^{\mathcal{D}},\ d = 0,\ldots,D;$$

2. Simulate L independent copies of (X_i, β_i):
$(X_i(\lambda), \beta_i(\lambda);\ i = 1,\ldots,n),\ \lambda = 1,\ldots,L.$
3. Solve the least squares problem

$$\begin{aligned} a_d(i;z) &= \arg\min_{a\in\mathbb{R}^K} \frac{1}{L}\sum_{\lambda=1}^{L}\left(\beta_{i+1}^{(d)}(\lambda) z(X_{i+1}(\lambda)) - \psi_d(i,X_i(\lambda))a\right)^2 \\ &\approx \arg\min_{a\in\mathbb{R}^K} E\left[\left(\beta_{i+1}^{(d)} z(X_{i+1}) - \psi_d(i,X_i)a\right)^2\right]; \end{aligned}$$

4. Define, as approximation for $E[\beta_{i+1}z(X_{i+1})|X_i = x]$,

$$\hat{E}[\beta_{i+1}z(X_{i+1})|X_i = x] = (\psi_0(i,x)a_0(i;z),\ldots,\psi_D(i,x)a_D(i;z))^\top .$$

Back to the BSDE case

- Lemor, Gobet Warin (2006): Apply least-squares MC with indicator functions of hypercubes as basis functions.
- Problem: The number of regression paths Λ, which is required to make the whole scheme converge at the order of $n^{-1/2}$ increases like $n^{2\mathcal{D}+3}$.
- Infeasible for e.g. $\mathcal{D} = 5$.
- Gobet and Turkedjiev (2016) show that the complexity and the number of paths which need be stored at the same time can be reduced by analyzing an explicit variant of the forward scheme of Bender and Denk (2007). E.g. for $\mathcal{D} = 5$, $n^{4.5}$ samples need be stored at the same time. ($40^{4.5} \approx 16$ Mio.)
- There is a curse of dimensionality, when a fine time grid is necessary or the dimension of X is large.

The primal-dual approach

- Idea: Assume that approximations $\hat{z}(i,x)$ of $z(i,x)$ are pre-computed by an algorithm of one's choice, e.g. the above approximate dynamic programming approach. Apply $\hat{z}(i,x)$ in order to compute a confidence interval for the quantity of interest Y_0 via Monte Carlo.
- Popular in Bermudan option pricing (Rogers, 2002; Haugh, Kogan, 2004; Andersen, Broadie, 2004).

The Bermudan option case

- The dynamic program for Bermudan option pricing (i.e. optimal stoping problem) is

$$Y_i = \max\{S_i, E_i[Y_{i+1}]\}, \quad Y_n = S_n,$$

i.e. $D = 0$, $\beta_i^{(0)} = 1$, $F(i, z) = \max\{S_i, z\}$ for some adapted process S (discounted payoff).

- Primal problem: Y_i is the value process of the following maximization problem (optimal stopping):

$$Y_i = \operatorname*{esssup}_{\tau} E_i[S_\tau],$$

where τ runs over the set of $\{i, \dots, n\}$-valued stopping times. An optimal stopping time is given by

$$\tau_i^* = \inf\{j \geq i;\ S_j \geq E_j[Y_{j+1}]\} \wedge n$$

The Bermudan option case

- Dual problem: Y_i is the value process of the following minimization problem (information relaxation dual):

$$Y_i = \operatorname*{essinf}_{M} E_i[\max_{j=i,\dots,n}(S_j - M_j)] + M(i)$$

where M runs over the set of martingales (Rogers, 2002; Haugh, Kogan, 2004). The martingale part M^* of the Doob decomposition of Y is optimal, even in a pathwise sense:

$$Y_i = \max_{j=i,\dots,n}(S_j - M_j^*) + M_i^*$$

The Bermudan option case

- Andersen/Broadie-algorithm (2004); very roughly speaking: Apply the approximation $\hat{z}(i,x)$ to $z(i,x) = E[Y_{i+1}|X_i = x]$ in order to construct a stopping time $\hat{\tau}_0$ and a martingale $\hat{M}$, which are 'close' to the optimal ones τ_0^* and M^*.
- Estimate
$$E[S_{\hat{\tau}_0}]$$
by plain Monte Carlo to get a lower confidence bound for Y_0 and
$$E[\max_{j=0,\ldots,n}(S_j - \hat{M}_j)] + E[\hat{M}_0]$$
to get an upper confidence bound.
- In practice: Use only a moderate effort to pre-calculate $\hat{z}$, e.g. least-squares MC with just a couple of well-chosen basis functions and a few thousand regression paths. Check whether the confidence interval by the primal-dual approach is sufficiently tight for the application under consideration.

The Bermudan option case – proof of the dual representation

- For any martingale M with $M_0 = 0$ it holds, by optional sampling, that

$$Y_0 = \sup_{\tau} E[S_\tau - M_\tau] \leq E[\max_{i=0,\ldots,n}(S_i - M_i)]$$

- Suppose

$$Y_i = Y_0 + M_i^* - A_i^*$$

is the Doob decomposition of Y_i. Then, by the supermartingale property of Y,

$$\begin{aligned} Y_0 &= \max_{i=0,\ldots,n}(Y_0 - A_i^*) \\ &= \max_{i=0,\ldots,n}(Y_i - M_i^*) \geq \max_{i=0,\ldots,n}(S_i - M_i^*) \end{aligned}$$

The Bermudan option case – pathwise dynamic programming

- Alternative approach to the dual representation.
- Define for a martingale M:

$$\theta_i^{up} := \max_{j=i,\dots,n}(S_j - M_j) + M_i := \max\{S_i, \theta_{i+1}^{up} - (M_{i+1} - M_i)\}.$$

- Then for $Y_i^{up} := E_i[\theta_i^{up}]$ by convexity of $z \mapsto \max\{S_i, z\}$

$$Y_i^{up} \geq \max\{S_i, E_i[\theta_{i+1}^{up} - (M_{i+1} - M_i)]\} = \max\{S_i, E_i[Y_{i+1}^{up}]\}.$$

- Hence, Y_i^{up} is a 'supersolution' for the dynamic program, and by backward induction and monotonicity of $z \mapsto \max\{S_i, z\}$,

$$Y_i^{up} \geq \max\{S_i, E_i[Y_{i+1}^{up}]\} \geq \max\{S_i, E_i[Y_{i+1}]\} = Y_i$$

Back to the general setting

- Dynamic program:

$$Y_i = F(i, E_i[\beta_{i+1} Y_{i+1}]), \quad 0 \le i \le n-1; \quad Y_n = \xi$$

- Adapted processes Y_i^{up} (resp. Y_i^{low}) are called supersolution (resp. subsolution) to the above dynamic program, if

$$Y_i^{up} \ge F(i, E_i[\beta_{i+1} Y_{i+1}^{up}]), \quad 0 \le i \le n-1; \quad Y_n^{up} \ge \xi;$$

  and with '$\ge$' replaced by '$\le$' for the subsolution.
- (Conv) The map $z \mapsto F(i, z)$ is convex.

Pathwise dynamic programming: The convex case

- Suppose (Conv). Given a $\mathbb{R}^{1+D}$-valued martingale M define $\theta_i^{up} := \theta_i^{up}(M)$ via the pathwise dynamic program

$$\theta_i^{up} = F(i, \beta_{i+1}\theta_{i+1}^{up} - (M_{i+1} - M_i)), \quad 0 \leq i \leq n-1; \quad \theta_n^{up} = \xi;$$

- Then, by (Conv),

$$E_i[\theta_i^{up}] \geq F(i, E_i[\beta_{i+1} E_{i+1}[\theta_{i+1}^{up}]])$$

and hence $Y_i^{up} = E_i[\theta_i^{up}]$ is a supersolution.

- Important: Calculating Y_0^{up} only requires the evaluation of one expectation (and not of nested ones), but of course also the choice of a martingale M.

Pathwise dynamic programming: The convex case

- Now take M^* as the martingale part of βY, i.e.

$$M^*_{i+1} - M^*_i = \beta_{i+1} Y_{i+1} - E_i[\beta_{i+1} Y_{i+1}].$$

- Then, P-almost surely,

$$\theta_i^{*,up} := \theta_i^{up}(M^*) = Y_i,$$

because by backward induction on i,

$$\begin{aligned} \theta_i^{*,up} &= F(i, \beta_{i+1}\theta_{i+1}^{up,*} - (\beta_{i+1} Y_{i+1} - E_i[\beta_{i+1} Y_{i+1}])) \\ &= F(i, E_i[\beta_{i+1} Y_{i+1}])) = Y_i \end{aligned}$$

- Question: Under which assumptions does a comparison principle hold, i.e. are supersolution above the solution and subsolutions below the solution?

Pathwise dynamic programming: The convex case

Theorem

Suppose (Conv). Then, the following assertions are equivalent:

- *(Comp) For any subsolution Y_i^{low} and any supersolution Y_i^{up}: $Y_i^{up} \geq Y_i^{low}$ P-a.s.*
- *(Mono) If $y, \tilde{y}$ are two integrable random variables such that $y \geq \tilde{y}$ P-a.s., then*

$$F(i, E_i[\beta_{i+1} y]) \geq F(i, E_i[\beta_{i+1}\tilde{y}]) \; P\text{-a.s.}$$

In this case

$$Y_i = \operatorname*{essinf}_M E_i[\theta_i^{up}(M)],$$

where M runs over the set of $\mathbb{R}^{1+D}$-valued martingales. Moreover, the Doob martingale of $\beta_i Y_i$ is a minimizer (and is 'surely optimal').

How restrictive is (Mono) ?

- Back to the BSDE case: Denote by $[\Delta W_i]$, the Brownian increments $W_{t_{i+1}} - W_{t_i}$ truncated at c such that $c + h \leq L^{-1}$, where L denotes the Lipschitz constant of f. Then,

$$F(i, E_i[\beta_{i+1} y]) = E_i[y] + f(t_i, E_i[y], h^{-1} E_i[[\Delta W_i] y])h.$$

Hence, for $y \geq \tilde{y}$

$$\begin{aligned} & F(i, E_i[\beta_{i+1} y]) - F(i, E_i[\beta_{i+1} \tilde{y}]) \\ \geq \quad & E_i[y - \tilde{y}] - Lh(E_i[y - \tilde{y}] + h^{-1} E_i[|[\Delta W_i]|(y - \tilde{y})]) \\ \geq \quad & (1 - Lh - cL) E_i[(y - \tilde{y})] \geq 0. \end{aligned}$$

- The truncation is inessential for the behaviour of the time discretization scheme. Hence (Mono) is harmless in the Lipschitz BSDE case.

Pathwise dynamic programming: The convex case

- Question: How to construct subsolutions?
- Define the convex conjugate of F by

$$F^{\#}(i,\rho) = \sup_{z \in \mathbb{R}^{1+D}} \left(\rho^\top z - F(i,z) \right).$$

- An adapted process ρ is said to be an admissible control, if

$$\sum_{j=0}^{n-1} E[|F^{\#}(j,\rho_j)|] < \infty.$$

- Given an admissible control consider the pathwise linearization $\theta^{low} = \theta^{low}(\rho)$

$$\theta_i^{low} = \rho_i^\top \beta_{i+1} \theta_{i+1}^{low} - F^{\#}(i,\rho_i), \quad 0 \le i \le n-1; \quad \theta_n^{low} = \xi;$$

and define $Y_i^{low} = E_i[\theta_i^{low}]$.

Pathwise dynamic programming: The convex case

- Then, by adaptedness of ρ and noting that $\rho \in [-L, L]^{1+D}$.

$$\begin{aligned} Y_i^{low} &= \rho_i^\top E_i[\beta_{i+1}\theta_{i+1}^{low}] - F^\#(i,\rho_i) \\ &= \rho_i^\top E_i[\beta_{i+1} Y_{i+1}^{low}] - F^\#(i,\rho_i) \le F(i, E_i[\beta_{i+1} Y_{i+1}^{low}]), \end{aligned}$$

  because $F^{\#\#} = F$ by convexity and since F is defined on the whole $\mathbb{R}^{1+D}$.

- Applying measurable selection of a subgradient (Cheridito, Kupper, Vogelpoth, 2015), there is an admissible control ρ^*, which satisfies

$$(\rho_i^*)^\top E_i[\beta_{i+1} Y_{i+1}] - F^\#(i,\rho_i^*) = F(i, E_i[\beta_{i+1} Y_{i+1}]),$$

  and hence by induction turns the above inequality into equality.

Pathwise dynamic programming: The convex case

Theorem

Suppose (Conv) and (Mono). Then

$$Y_i = \operatorname*{esssup}_{\rho} E_i[\theta_i^{low}(\rho)],$$

where ρ runs over the set of admissible controls. Any admissible ρ^* satisfying

$$(\rho_i^*)^\top E_i[\beta_{i+1} Y_{i+1}] - F^\#(i, \rho_i^*) = F(i, E_i[\beta_{i+1} Y_{i+1}])$$

is optimal and such maximizer exists.

Information relaxation

- Solving the linear program for $\theta_0^{low}(\rho)$ explicitly yields the primal problem

$$Y_0 = \sup_{(\rho_j)_{j=0,\dots,n-1}} E[\xi \prod_{k=0}^{n-1} \rho_k^\top \beta_{k+1} - \sum_{i=0}^{n-1} F^\#(i,\rho_i) \prod_{k=0}^{i-1} \rho_k^\top \beta_{k+1}],$$

where the sup runs over admissible controls ρ.

- The information relaxation dual due to Brown, Smith, Sun (2010), cp. also Rogers (2007), states that

$$Y_0 = \inf_{\mathfrak{p}} E[\sup_{(r_i)_{i=0,\dots,n-1}} \Big(\xi \prod_{k=0}^{n-1} r_k^\top \beta_{k+1} - \sum_{i=0}^{n-1} F^\#(i,r_i) \prod_{k=0}^{i-1} r_k^\top \beta_{k+1} - \mathfrak{p}((r_i)_i)\Big)],$$

where $\mathfrak{p}$ in general runs over a huge class of 'penalties'.

- Our results show: one can restrict to 'martingale penalties':

$$\mathfrak{p}(r) = \sum_{i=0}^{n-1} r_i^\top (M_{i+1} - M_i) \prod_{k=0}^{i-1} r_k^\top \beta_{k+1}.$$

Monte Carlo implementation

- Step 1: Approximate dynamic programming: Pre-compute an approximation $\hat{z}(i,x)$ of

  $$z(i,x) = E[\beta_{i+1}Y_{i+1}|X_i = x] = E[\beta_{i+1}F(i{+}1, z_{i+1}(X_{i+1}))|X_i = x],$$

  initiated at $z(n,x) = 0$ with the convention $F(n,0) := \xi$. E.g.:

  $$\hat{z}(i,x) = \hat{E}[\beta_{i+1}F(i+1, \hat{z}_{i+1}(X_{i+1})|X_i = x], \quad \hat{z}(n,x) = 0,$$

  where $\hat{E}[\cdot|X_i = x]$ approximates conditional expectation operator. Let $\hat{Y}_0 := F(0, z(0, X_0))$.
- We apply least-squares Monte Carlo, i.e. we have to choose basis functions and the number of regression paths.
- In our implementation we choose basis functions which have a certain martingale property. Then, the martingale basis variant of least-squares MC (Bender, Steiner, 2012; cp. also Glasserman, Yu, 2004) can be run in order to reduce the variance.

Monte Carlo implementation

- Step 2: Lower confidence bound: Given $\hat{z}(i,x)$ there is a $\hat{\rho}_i(x)$ which solves
$$\hat{\rho}_i(x)^\top \hat{z}(i,x) - F^{\#}(i, \hat{\rho}_i(x)) = F(i, \hat{z}(i,x))$$
such that $\hat{\rho}_i(X_i)$ is admissible.
- Simulate Λ^{out} 'outer' sample paths $(\beta_i(\lambda), X_i(\lambda))$ of (β_i, X_i) (independent of whatever paths might have been used to compute $\hat{z}$) and define
$$\hat{\theta}_i^{low}(\lambda) = \hat{\rho}_i(X_i(\lambda))^\top \beta_{i+1}(\lambda)\theta_{i+1}^{low}(\lambda) - F^{\#}(i, \hat{\rho}_i(X_i(\lambda)))$$
initiated at the terminal condition of the dynamic program.
- Then, the plain MC estimator
$$\frac{1}{\Lambda^{out}} \sum_{\lambda=1}^{\Lambda^{out}} \hat{\theta}_i^{low}(\lambda)$$
is biased downwards for Y_0 and an asymptotic confidence bound can be constructed in the usual way.

Monte Carlo implementation

- Step 3: Upper confidence bound: Given $\hat{z}(i,x)$ define along each outer paths

$$\begin{aligned}\Delta \hat{M}_i(\lambda) &= \beta_{i+1}(\lambda) F(i+1, \hat{z}(i+1, X_{i+1}(\lambda))) \\ &\quad - E[\beta_{i+1} F(i+1, \hat{z}(i+1, X_{i+1}))|X_i = X_i(\lambda)].\end{aligned}$$

  If the conditional expectation is not available in closed form, replace it by an unbiased estimator, e.g. one layer of nested simulation with Λ^{in} 'inner' sample paths.

- Define

$$\hat{\theta}_i^{up}(\lambda) = F(i, \beta_{i+1}(\lambda)\hat{\theta}_{i+1}^{up}(\lambda) - \Delta \hat{M}_i(\lambda))$$

  initiated at the the terminal condition of the dynamic program and proceed analogously to the 'lower bound' in order to construct and estimator with bias upwards and an upper confidence bound.

A numerical example

- Given a 5-dimensional Brownian motion W consider the BSDE

$$
\begin{aligned}
dX_{d,t} &= \mu X_{d,t}dt + \sigma X_{d,t}dW_{d,t}, \quad X_{d,0} = x_0 \\
dY_t &= -\frac{R}{\sigma}\left(\sum_{d=1}^{5} Z_{d,t}\right)_+ dt + Z_t dW_t, \quad Y_T = k(\max_{d=1,\dots,5} X_{d,T})
\end{aligned}
$$

  It yields the *g*-expectation

$$
Y_0 = \operatorname*{esssup}_b E\left[k\left(\max_{d=1,\dots,5} x_0 \exp\left\{\sigma W_{d,T} + \int_0^T (b_s - \frac{\sigma^2}{2})ds\right\}\right)\right],
$$

  where b runs over the adapted processes with values in $[\mu, \mu + R]$.

- Choice of parameters: $k(x) = (x - 95)_+ - 2(x - 115)_+$

$$
T = 1/4,\ \mu = 0.01,\ R = 0.05,\ \sigma = 0.2,\ x_0 = 100.
$$

A numerical example

- Computation of $\tilde{z}$: Least-squares MC with martingale basis enhancement, 100 regression paths, and basis functions

$$\psi_{0,1}(i,x) = 1, \quad \psi_{0,2}(i,x) = E[k(\max_{d=1,\ldots,5} X_{d,T})|X_{t_i} = x]$$

$$\psi_{d,1}(i,x) = x_d \frac{d}{dx_d} E[k(\max_{d=1,\ldots,5} X_{d,T})|X_{t_i} = x], \; d = 1,\ldots,D.$$

- Computation of the confidence bounds: 10.000 outer paths, 100 inner paths; use of various control variates.
- Numerical results (standard deviations in brackets):

n	40	80	120	160
$\hat{Y}_0^{low}$	13.936 (0.003)	13.935 (0.003)	13.941 (0.003)	13.942 (0.003)
$\hat{Y}_0^{up}$	13.976 (0.003)	14.001 (0.003)	14.033 (0.003)	14.061 (0.003)

Uncertain volatility

- Consider the *G*-expectation

$$\operatorname*{esssup}_{\sigma} E\left[k\left(x_0 \exp\left\{\int_0^T \sigma(s)dW_s - \int_0^T \frac{\sigma^2(s)}{2}ds\right\}\right)\right],$$

where σ runs over the adapted processes with values in $[\underline{\sigma}, \overline{\sigma}]$.

- This G-expectation can be described by a second order BSDE, which after a time discretization becomes (Fahim, Touzi, Warin, 2011; Guyon, Henry-Labordere, 2011)

$$Y_i = E_i[Y_{i+1}] + \max_{\rho \in [(\underline{\sigma}/\hat{\sigma}-1)/2, (\overline{\sigma}/\hat{\sigma}-1)/2]} \rho E_i[(\Delta W_{i+1}^2/h - \rho \Delta W_{i+1} - 1)Y_{i+1}]$$

with terminal condition $Y_n = k(x_0 \exp\{\hat{\sigma} W_T - \hat{\sigma}^2 T/2\})$.

- (Mono) is violated in the popular numerical test case: $\underline{\sigma} = 0.1$, $\overline{\sigma} = 0.2$ for any choice of $\hat{\sigma}$.

Beyond the monotonicity condition

- Suppose (Conv), but that (Mono) is violated and thus the comparison principle is not in force.
- Given an $\mathbb{R}^{1+D}$-valued martingale M with increments ΔM and an admissible control ρ define $\theta^{up} = \theta^{up}(M, \rho)$ and $\theta^{low} = \theta^{low}(M, \rho)$ via

$$\begin{aligned} \theta_i^{up} &= \max\{F(i, \beta_{i+1}\theta_{i+1}^{up} - \Delta M_{i+1}),\ F(i, \beta_{i+1}\theta_{i+1}^{low} - \Delta M_{i+1})\} \\ \theta_i^{low} &= (\rho_i^\top \beta_{i+1})_+ \theta_{i+1}^{low} - (\rho_i^\top \beta_{i+1})_- \theta_{i+1}^{up} - \rho_i^\top \Delta M_{i+1} - F^\#(i, \rho_i), \end{aligned}$$

  initiated at $\theta_n^{up} = \theta_n^{low} = \xi$.
- Then, $E_i[\theta_i^{up}(M, \rho)]$ is a supersolution, $E_i[\theta_i^{low}(M, \rho)]$ is a subsolution and

$$E_i[\theta_i^{low}(M, \rho)] \le Y_i \le E_i[\theta_i^{up}(M, \rho)].$$

Beyond the monotonicity condition

Theorem

Suppose (Conv). Then

$$Y_i = \operatorname*{esssup}_{M,\rho} E_i[\theta_i^{low}(M,\rho)] = \operatorname*{essinf}_{M,\rho} E_i[\theta_i^{up}(M,\rho)]$$

where ρ runs over the set of admissible controls and M over the set of $\mathbb{R}^{1+D}$-valued martingale M . Optimizers (M^*, ρ^*) for both problems are given by the Doob martingale M^* of βY and any admissible ρ^* satisfying

$$(\rho_i^*)^\top E_i[\beta_{i+1} Y_{i+1}] - F^\#(i, \rho_i^*) = F(i, E_i[\beta_{i+1} Y_{i+1}]).$$

These optimizers are 'surely' optimal, i.e. P-almost surely

$$Y_i = \theta_i^{low}(M^*, \rho^*) = \theta_i^{up}(M^*, \rho^*)$$

A numerical example

- Recall the pricing problem under uncertain volatility

$$\sup_{\sigma} E\left[g\left(S_0 \exp\left\{\int_0^T \sigma_s dW_s - \int_0^T \frac{\sigma_s^2}{2} ds\right\}\right)\right],$$

  where σ runs over the progressively measurable processes with values in $[\underline{\sigma}, \overline{\sigma}]$.

- Time discretization with reference volatility $\hat{\sigma}$

$$Y_i = E_i[Y_{i+1}] + \frac{1}{2} \max_{\rho \in \{\underline{\sigma}, \overline{\sigma}\}} ((\frac{\rho^2}{\hat{\sigma}^2} - 1) E_i[Y_{i+1}(\frac{\Delta W_{i+1}^2}{h} - \hat{\sigma}\Delta W_{i+1} - 1)])$$

$$Y_n = g(S_0 \exp\{\hat{\sigma} W_T - \hat{\sigma}^2 T/2\}),$$

- Least-squares Monte Carlo for the approximate dynamic program suffers from large variances due to the gamma weight, see the discussion in Alanko and Avellaneda (2013).

A numerical example

- Choice of the parameters:

$$\underline{\sigma} = 0.1, \quad \overline{\sigma} = 0.2, \quad T = 1, \quad S_0 = 100$$

- Call spread option:

$$g(x) = (x - 90)_+ - (x - 110)_+$$

- Approximate dynamic programming: martingale basis variant of least-squares MC, 100.000 regression paths, basis functions: $1, x$, and the prices of call options with 160 different strikes between 20.5 and 230.5.
- Computation of the confidence bounds: 100.000 outer paths, martingale available in closed form.

A numerical example

- $\hat{\sigma} = 0.15$:

n	$\hat{Y}_0$	$\hat{Y}_0^{low}$	$\hat{Y}_0^{up}$	$\hat{Y}_0^{LS}$
5	10.8153	10.8153 (0.0001)	10.8167 (0.0001)	10.9391
15	11.0684	11.0683 (0.0001)	11.0728 (0.0001)	11.6781
25	11.1241	11.1237 (0.0001)	11.1379 (0.0008)	12.3851
35	11.1479	11.1462 (0.0002)	11.2058 (0.0047)	13.3329

- $\hat{\sigma} = 0.2/\sqrt{3} \approx 0.12$:

n	$\hat{Y}_0$	$\hat{Y}_0^{low}$	$\hat{Y}_0^{up}$	$\hat{Y}_0^{LS}$
3	10.8553	10.8550 (0.0001)	10.8591 (0.0001)	10.9440
9	11.1054	11.1060 (0.0005)	11.1116 (0.0005)	11.7690
15	11.1484	11.1486 (0.0002)	11.1728 (0.0056)	13.1633
21	11.1666	11.1666 (0.0003)	11.2764 (0.0173)	13.7743

- True option price in continuous time (Vanden, 2006): 11.2046

Beyond the convexity condition

- Suppose (Mono), which ensures the comparison pinciple, even if (Conv) is violated. Denote by L a Lipschitz constant of F.
- Fix an integrable process $\tilde{Z}_i$, which we think of as an approximation of $Z_i = E_i[\beta_{i+1} Y_{i+1}]$
- Consider the auxiliary convex dynamic program

$$Y_i^{\tilde{Z}} = F(i, \tilde{Z}_i) + L \left| \tilde{Z}_i - E_i[\beta_{i+1} Y_{i+1}^{\tilde{Z}}] \right|, \quad Y_n^{\tilde{Z}} = \xi.$$

  Then $Y^{\tilde{Z}}$ is a supersolution to the original dynamic program for any $\tilde{Z}$ and $Y^Z = Y$.
- Hence

$$Y_i = \operatorname*{essinf}_{\tilde{Z}} \operatorname*{essinf}_{M} E_i[\theta_i^{up,\tilde{Z}}(M)].$$

- For the 'lower bound' one considers the auxiliary concave dynamic program

$$\underline{Y}_i^{\tilde{Z}} = F(i, \tilde{Z}_i) - L \left| \tilde{Z}_i - E_i[\beta_{i+1} \underline{Y}_{i+1}^{\tilde{Z}}] \right|, \quad \underline{Y}_n^{\tilde{Z}} = \xi.$$

A numerical example

- Adjusting option pricing problems for default risk can be incorporated in the BSDE setting with $D = 0$, $\beta = 1$ by considering the generator

$$f(i, y) = -(1 - \delta)Q(y)y - Ry,$$

where R is the default free rate, δ is the recovery rate and the default itensity depends on the option value through the function Q, see Duffie, Schroder, Skiadas (1996).

- We choose Q piecewise linear in a way that f is neither convex nor concave.
- Option payoff $Y_n = \min_{d=1,\dots,5} X_T^d$, where X^d are independent, identically distributed Black-Scholes stocks with $X_0 = 100$, $\sigma = 0.2$, $\mu = R = 0.05$, and $T = 1$.

A numerical example

- Approximate dynamic program: Least-squares MC with 2 basis functions 1, $E[\min_{d=1,\dots,5} X_T^d | X_{t_i} = x]$, 10^5 regression paths.
- Confidence bounds: 4.000 outer paths and 1.000 inner paths.

$\delta \setminus n$	40		80		120	
0	71.6551 (0.0071)	71.8589 (0.0068)	71.6774 (0.0072)	71.8828 (0.0068)	71.6664 (0.0070)	71.8656 (0.0068)
$\frac{1}{3}$	74.1023 (0.0062)	74.2241 (0.0060)	74.1010 (0.0065)	74.2225 (0.0062)	74.1032 (0.0062)	74.2229 (0.0061)
$\frac{2}{3}$	76.3335 (0.0057)	76.3865 (0.0057)	76.3364 (0.0057)	76.3886 (0.0057)	76.3416 (0.0059)	76.3943 (0.0058)

Conclusion

- Generalizing from the Bermudan option case, we presented a methodology which complements any numerical method for approximating discrete time BSDEs (and more general dynamic programs) with a confidence interval for the quantity of interest Y_0.
- In particular, a ‘cheap’ approximation method for the conditional expectations in the approximate dynamic program can be justified a-posteriori, if the confidence interval is sufficiently tight for the application under consideration.

Thank you for your attention

This talk was based on:

- Bender, Schweizer, Zhuo, A primal-dual algorithm for BSDEs. Mathematical Finance. Early view
- Bender, Gärtner, Schweizer, Pathwise dynamic programming. Preprint.

This research was supported by the Deutsche Forschungsgemeinschaft via the priority program 1324.